

Hortonworks Data Platform

Hadoop High Availability

(Dec 2, 2014)

Hortonworks Data Platform: Hadoop High Availability

Copyright © 2012-2015 Hortonworks, Inc. Some rights reserved.

The Hortonworks Data Platform, powered by Apache Hadoop, is a massively scalable and 100% open source platform for storing, processing and analyzing large volumes of data. It is designed to deal with data from many sources and formats in a very quick, easy and cost-effective manner. The Hortonworks Data Platform consists of the essential set of Apache Hadoop projects including MapReduce, Hadoop Distributed File System (HDFS), HCatalog, Pig, Hive, HBase, Zookeeper and Ambari. Hortonworks is the major contributor of code and patches to many of these projects. These projects have been integrated and tested as part of the Hortonworks Data Platform release process and installation and configuration tools have also been included.

Unlike other providers of platforms built using Apache Hadoop, Hortonworks contributes 100% of our code back to the Apache Software Foundation. The Hortonworks Data Platform is Apache-licensed and completely open source. We sell only expert technical support, [training](#) and partner-enablement services. All of our technology is, and will remain free and open source.

Please visit the [Hortonworks Data Platform](#) page for more information on Hortonworks technology. For more information on Hortonworks services, please visit either the [Support](#) or [Training](#) page. Feel free to [Contact Us](#) directly to discuss your specific needs.



Except where otherwise noted, this document is licensed under
Creative Commons Attribution ShareAlike 3.0 License.
<http://creativecommons.org/licenses/by-sa/3.0/legalcode>

Table of Contents

1. High Availability for Hive Metastore	1
1.1. Use Cases and Failover Scenarios	1
1.2. Software Configuration	2
1.2.1. Install HDP	2
1.2.2. Update the Hive Metastore	3
1.2.3. Validate Configuration	3
2. Highly Available Reads with HBase	4
2.1. Introduction to HBase High Availability	5
2.2. Propagating Writes to Region Replicas	7
2.3. Timeline Consistency	9
2.4. Configuring HA Reads for HBase	11
2.5. Creating Highly-Available HBase Tables	13
2.6. Querying Secondary Regions	14
2.7. Monitoring Secondary Region Replicas	15
3. Namenode High Availability	16
3.1. Architecture	16
3.2. Hardware Resources	17
3.3. Deploy NameNode HA Cluster	18
3.3.1. 1. Configure NameNode HA Cluster	18
3.3.2. 2. Deploy NameNode HA Cluster	23
3.3.3. 3. Deploy Hue with an HA Cluster	26
3.3.4. 4. Deploy Oozie with an HA Cluster	27
3.4. Operating a NameNode HA Cluster	28
3.5. Configure and Deploy NameNode Automatic Failover	29
3.5.1. Prerequisites	30
3.5.2. Instructions	31
3.5.3. Configuring Oozie Failover	32
3.6. Appendix: Administrative Commands	33
4. Resource Manager High Availability	35
4.1. Hardware Resources	35
4.2. Deploy ResourceManager HA Cluster	35
4.2.1. Configure Manual or Automatic ResourceManager Failover	36
4.2.2. Deploy the ResourceManager HA Cluster	39
4.2.3. Minimum Settings for Automatic ResourceManager HA Configuration	40
4.2.4. Testing ResourceManager HA on a Single Node	41

1. High Availability for Hive Metastore

This document is intended for system administrators who need to configure the Hive Metastore service for High Availability. It contains the following sections:

- [Use Cases and Failover Scenarios](#)
- [Software Configuration](#)

The following subsequent chapters address high availability with related components:

- [Highly Available Reads with HBase](#)
- [Namenode High Availability](#)
- [Resource Manager High Availability](#)



Important

The relational database that backs the Hive metastore itself should also be made highly available using best practices defined for the database system in use.

1.1. Use Cases and Failover Scenarios

This section provides information on the use cases and failover scenarios for high availability (HA) in the Hive metastore.

Use Cases

The metastore HA solution is designed to handle metastore service failures. When a deployed metastore service fails, it can remain unavailable for a considerable time until the service is started again. To avoid such outages, deploy the metastore service in HA mode.

Deployment Scenarios

We recommend deploying the metastore service on multiple servers concurrently. Each Hive metastore client reads the configuration setting in the `hive.metastore.uris` property to get a list of metastore servers with which it can attempt to communicate:

```
<property>
  <name>hive.metastore.uris</name>
  <value>thrift://$Hive_Metastore_Server_Host_Machine_FQDN</value>
  <description>A comma separated list of metastore uris on which metastore
  service is running</description>
</property>
```

These metastore servers store their state in a MySQL HA cluster. See the MySQL documentation for information about configuring a MySQL HA cluster.

In the case of a secure cluster, each of the metastore servers will additionally need to have the following configuration property in its `hive-site.xml` file:

```
<property>
  <name>hive.cluster.delegation.token.store.class</name>
  <value>org.apache.hadoop.hive.thrift.DBTokenStore</value>
</property>
```

Failover Scenario

A Hive metastore client always uses the first URI to connect with the metastore server. In case the metastore server becomes unreachable, the client randomly picks a URI from the list and tries to connect with it.

1.2. Software Configuration

Complete the following tasks to configure a Hive HA solution:

1. [Install HDP](#)
2. [Update the Hive Metastore](#)
3. [Validate configuration](#)

1.2.1. Install HDP

Use the following instructions to install HDP on your cluster hardware. Ensure that you specify the virtual machine that you configured in the section, "[Use Cases and Failover Scenarios](#)" as your NameNode.

1. Download the Apache Ambari repository using the instructions provided in "[Ambari Repositories](#)" section of the *Ambari 1.7.0 Documentation Suite*.



Note

Do not start the Ambari server until you have configured the relevant templates as outlined in the following steps.

2. Edit the `<master-install-machine-for-Hive-Metastore>/etc/hive/conf.server/hive-site.xml` configuration file to add the following properties:

- Provide the URI for the client to contact the metastore server. The following property can have a comma separated list when your cluster has multiple Hive metastore servers.

```
<property>
  <name>hive.metastore.uris </name>
  <value>thrift://$Hive_Metastore_Server_Host_Machine_FQDN</value>
  <description>URI for client to contact metastore server</description>
</property>
```

- Configure the Hive cluster delegation token storage class.

```
<property>
  <name>hive.cluster.delegation.token.store.class</name>
  <value>org.apache.hadoop.hive.thrift.DBTokenStore</value>
</property>
```

- Complete the HDP installation.
- Continue the Ambari installation process using the instructions provided in "[Install, Configure and Deploy an HDP Cluster](#)" section of the *Ambari 1.7.0 Documentation Suite*.
- Complete the Ambari installation. Ensure that the installation is successful.

1.2.2. Update the Hive Metastore

HDP components configured for HA must use a NameService rather than a NameNode. Use the following instructions to update the Hive metastore to reference the NameService rather than a Name Node.



Note

Hadoop administrators also use the following procedure to update the Hive metastore with the new URI for a node in a Hadoop cluster. For example, administrators sometimes rename an existing node as their cluster grows.

1. Open a command prompt on the machine hosting the Hive metastore.
2. Set the `HIVE_CONF_DIR` environment variable:

```
export HIVE_CONF_DIR=/etc/hive/conf.server
```

3. Execute the following command to retrieve a list of URIs for the filesystem roots, including the location of the NameService:

```
hive --service metatool -listFSRoot
```

4. Execute the following command with the `-dryRun` option to test your configuration change before implementing it:

```
hive --service metatool -updateLocation <nameservice-uri> <namenode-uri> -dryRun
```

5. Execute the command again, this time without the `-dryRun` option:

```
hive --service metatool -updateLocation <nameservice-uri> <namenode-uri>
```

1.2.3. Validate Configuration

Test various failover scenarios to validate your configuration.

2. Highly Available Reads with HBase

HDP 2.2 enables HBase administrators to configure HBase clusters with read-only high availability (HA). This feature benefits HBase applications that require low-latency queries and can tolerate minimal staleness for read operations. Examples include queries on remote sensor data, distributed messaging, object stores, and user profile management.

High Availability for HBase features the following functionality:

- Data is safely protected in HDFS
- Failed nodes are automatically recovered
- No single point of failure
- All HBase API and region operations are supported, including scans, region split/merge, and META table support. The META table stores information about regions.

However, HBase administrators should carefully consider the following costs associated with using high availability features:

- Double or triple `MemStore` usage
- Increased `BlockCache` usage
- Increased network traffic for log replication
- Extra backup RPCs for secondary region replicas

HBase is a distributed key-value store designed for fast table scans and read operations at petabyte scale. Before configuring HA for HBase, you should understand the concepts in the following table.

HBase Concept	Description
Region	<p>A group of contiguous rows in an HBase table. Tables start with one region. Additional regions are added dynamically as the table grows. Regions can be spread across multiple hosts to balance workloads and recover quickly from failure.</p> <p>There are two types of regions: primary and secondary. A secondary region is a copy of a primary region that is replicated on a different Region Server.</p>
Region Server	<p>A Region Server serves data requests for one or more regions. A single region is serviced by only one Region Server, but a Region Server may serve multiple regions. When region replication is enabled, a Region Server can serve regions in primary and secondary mode concurrently.</p>
Column family	<p>A column family is a group of semantically related columns that are stored together.</p>
MemStore	<p><code>MemStore</code> is in-memory storage for a Region Server. Region Servers write files to HDFS after the <code>MemStore</code> reaches a configurable maximum value specified with the <code>hbase.hregion.memstore.flush.size</code> property in the <code>hbase-site.xml</code> configuration file.</p>

HBase Concept	Description
Write Ahead Log (WAL)	The WAL is a log file that records all changes to data until the data is successfully written to disk and the MemStore is flushed). This protects against data loss in the event of a failure before MemStore contents are written to disk.
Compaction	When operations stored in the MemStore are flushed to disk, HBase consolidates and merges many smaller files into fewer large files. This consolidation is called <i>compaction</i> , and it is usually very fast. However, if many Region Servers hit the data limit that is specified by the MemStore at the same time, HBase performance may degrade from the large number of simultaneous major compactions. Administrators can avoid this by manually splitting tables over time.

For information about configuring regions, see "HBase Cluster Capacity and Region Sizing" in the [System Administration Guides](#).

2.1. Introduction to HBase High Availability


Architecturally, HBase has had a strong consistency guarantee from the start. All reads and writes are routed through a single Region Server, which guarantees that all writes happen in order, and all reads access the most recently committed data.

However, because of this "single homing" of reads to a single location, if the server becomes unavailable, the regions of the table that are hosted in the Region Server become unavailable for some time until they are recovered. There are three phases in the region recovery process: detection, assignment, and recovery. Of these, the detection phase is usually the longest, currently 20 to 30 seconds, depending on the Zookeeper session timeout setting when the Region Server fails but the Zookeeper session is running. After that, data is recovered from the Write Ahead Log and assigns the region to a different server. Until the recovery is complete, during this time clients cannot read data from that region.


For some use cases, the data may be read-only or stale data is acceptable. However, for use cases where latency is not acceptable, HBase can be used. To achieve high availability for reads, HBase provides a feature called *region replication*. When region replication is used each region of a table can be replicated and opened in different Region Servers. By default, the region replication is set to 1, so only a single region replica is deployed and there will not be any changes from the original model. If region replication is set to 2 or more, then the master assigns replicas of the regions of the table. The Load Balancer ensures that the region replicas are not co-hosted in the same Region Servers and in the same rack if possible.

All of the replicas for a single region have a unique replica ID, starting with 0. The region replica with replica ID = 0 is called the "primary region." The others are called "secondary region replicas," or "secondaries." Only the primary region can accept writes from the client, and the primary always contains the latest changes. Since all writes must go through the primary region, writes are not highly available because they might be blocked for some time if the region becomes unavailable.

For example, in the following image Region Server 1 is responsible for responding to queries and scans for keys 10 through 40. If Region Server 1 crashes, the region holding keys 10-40 is unavailable for a short time until the region recovers.



HA provides a way to access keys 10-40 even if Region Server 1 is not available, by hosting replicas of the region and assigning the region replicas to other Region Servers as backups. In the following image, Region Server 2 hosts secondary region replicas for keys 10-20, and Region Server 3 hosts the secondary region replica for keys 20-40. Region Server 2 also hosts the secondary region replica for keys 80-100. There are no separate Region Server processes for secondary replicas. Rather, Region Servers can serve regions in primary or secondary mode. When Region Server 2 services queries and scans for keys 10-20, it acts in secondary mode.



Note

Regions acting in secondary mode are also known as Secondary Region Replicas. However, there is no separate Region Server process. A region in secondary mode can read but cannot write data. In addition, the data it returns may be stale, as described in the following section.

Timeline and Strong Data Consistency

HBase guarantees timeline consistency for all data served from Region Servers in secondary mode. This means that all HBase clients see the same data in the same order, but that data may be slightly stale. Only the primary Region Server is guaranteed to have the latest data. Timeline consistency simplifies the programming logic for complex HBase queries and provides lower latency than quorum-based consistency.

In contrast, strong data consistency means that the latest data is always served. However, strong data consistency can greatly increase latency in the case of a Region Server failure, because only the primary Region Server is guaranteed to have the latest data. The HBase API allows application developers to specify the data consistency that is required for a query.



Note

The HBase API contains a method called `Result.isStale()`, which specifies whether data returned in secondary mode can be stale, which means that the data has not been updated with the latest write operation to the primary Region Server.

2.2. Propagating Writes to Region Replicas

Writes are written only to the primary region replica.



Note

By default, HBase tables do not use high availability features. After configuring your cluster for high availability, designate tables as HA by setting region replication to a value greater than 1 at table creation time. For more information, see [Creating Highly-Available HBase Tables](#).

For read-only tables, you do not need to use any of the following methods. Disabling and enabling the table should make the data available in all region replicas.

The following two mechanisms are used to propagate writes from the primary replica to secondary replicas.

StoreFile Refresher

The first mechanism is the StoreFile refresher, which was introduced in Phase 1 (Apache HBase 1.0.0 and HDP 2.1).

StoreFile refresher is a thread per Region Server, which runs periodically, and does a refresh operation for the store files of the primary region for the secondary region replicas. If enabled, the refresher ensures that the secondary region replicas see the new flushed, compacted, or bulk loaded files from the primary region in a timely manner. However, this means that only flushed data can be read back from the secondary region replicas, and after the refresher runs, making the secondaries lag behind the primary for a longer time.

To enable this feature, configure `hbase.regionserver.storefile.refresh.period` to a value greater than

zero, and set `hbase.regionserver.storefile.refresh.all` to true. For more information about these properties, see [Configuring HA Reads for HBase](#).

Async WAL Replication

The second mechanism for propagating writes to secondaries is done via the Async WAL Replication feature. This feature is only available in HA Phase 2 (starting with HDP 2.2).

Async WAL replication works similarly to HBase's multi-datacenter replication, but the data from a region is replicated to its secondary regions. Each secondary replica always receives writes in the same order that the primary region committed them. In some sense, this design can be thought of as "in-cluster replication"; instead of replicating to a different datacenter, the data goes to secondary regions. This process keeps the secondary region's in-memory state up to date. Data files are shared between the primary region and the other replicas, so there is no extra storage overhead. However, secondary regions will have recent non-flushed data in their MemStores, which increases memory overhead. The primary region writes flush, compaction, and bulk load events to its WAL as well, which are also replicated through WAL replication to secondaries. When secondary replicas detect a flush/compaction or bulk load event, they replay the event to pick up the new files and drop the old ones.

Committing writes in the same order as in the primary region ensures that the secondaries won't diverge from the primary region's data, but because the log replication is asynchronous, the data might still be stale in secondary regions. Because this feature works as a replication endpoint, performance and latency characteristics should be similar to inter-cluster replication.

Async WAL Replication is disabled by default in HA Phase 2. To enable this feature, set `hbase.region.replica.replication.enabled` to true and set `hbase.regionserver.storefile.refresh.all` to false. For more information about these properties, see [Creating Highly-Available HBase Tables](#).

When you create a table with high availability enabled, the Async WAL Replication feature adds a new replication peer named `region_replica_replication`.

Once enabled, to disable this feature you must perform the following steps:

1. Set `hbase.region.replica.replication.enabled` property to **false** in `hbase-site.xml`.
2. In your cluster, disable the replication peer named `region_replica_replication`, using `hbase` shell or `ReplicationAdmin` class:

```
hbase> disable_peer 'region_replica_replication'
```

Store File TTL

In phase 1 and 2 of the write propagation approaches described previously, store files for the primary replica will be opened in secondaries independent of the primary region. Thus, for files that the primary region compacted and archived, the secondaries might still refer to these files to read data.

Both features use `HFileLinks` to refer to files, but the file might be deleted prematurely. To prevent I/O exceptions for requests to replicas, set the configuration property `hbase.master.hfilecleaner.ttl` to a sufficient time range such as 1 hour.

Region Replication for the Region of the META Table

Currently, Async Write Ahead Log (WAL) Replication is not done for the write ahead log of the META table. The secondary replicas of the META table refresh themselves from the persistent store files. To ensure that the META store files are refreshed, set the `hbase.regionserver.storefile.refresh.period` property to a non-zero value.



Note

Setting the configuration property `hbase.regionserver.storefile.refresh.all` to `false` only refreshes the store files of the META table's region.

2.3. Timeline Consistency

With timeline consistency, HBase introduces a consistency definition that can be provided per get or scan read operation:

```
public enum Consistency {  
    STRONG,  
    TIMELINE  
}
```

`Consistency.STRONG` is the default consistency model provided by HBase. If a table has region replication set to **1**, or has region replicas but the reads are done with time consistency enabled, the read is always performed by the primary regions. This preserves previous behavior and the client receives the latest data.

If a read is performed with `Consistency.TIMELINE`, then the read RPC will be sent to the primary Region Server first. After a short interval, such as the default setting of 10 milliseconds for the `hbase.client.primaryCallTimeout.get` property, the parallel RPC for secondary region replicas is sent if the primary region replica does not respond. HBase returns the result from the RPC that finishes first. If the response is from the primary region replica, the data is current. You can use the `Result.isStale()` API to determine the state of the resulting data:

- If the result is from a primary region, `Result.isStale()` is set to **false**.
- If the result is from a secondary region, `Result.isStale()` is set to **true**.

`TIMELINE` consistency as implemented by HBase differs from pure eventual consistency in the following respects:

- Single homed and ordered updates: Whether region replication is enabled or not, on the write side, there is still only one defined replica, the primary, that can accept writes. This replica is responsible for ordering the edits and preventing conflicts. This guarantees that two different writes are not committed at the same time by different replicas, resulting in divergent data. With this approach, there is no need to do read-repair or last-timestamp-wins types of conflict resolution.
- The secondary replicas also apply edits in the order that the primary committed them. Thus the secondaries contain a snapshot of the primary data at any point in time. This is

similar to RDBMS replications and HBase multi-datacenter replication, but takes place in a single cluster.

- On the read side, the client can detect whether the read is coming from up-to-date data or stale data. Also, the client can issue reads with different consistency requirements on a per-operation basis to ensure its own semantic guarantees.
- The client might still read stale data if it receives data from one secondary replica first, followed by reads from another secondary replica. There is no stickiness to region replicas, nor is there a transaction ID-based guarantee. If required, this can be implemented later.

Memory Accounting

Secondary region replicas refer to data files in the primary region replica, but they have their own `MemStores` in HA Phase 2 and use block cache as well. However, secondary region replicas cannot flush data when there is memory pressure for their `MemStores`. They can only free up `MemStore` memory when the primary region does a flush and the flush is replicated to the secondary.

Because a Region Server can host primary replicas for some regions and secondaries for others, secondary replicas might generate extra flushes to primary regions in the same host. In extreme situations, there might be no memory for new writes from the primary, by way of write ahead log (WAL) replication.

To resolve this situation, the secondary replica is allowed to do a *store file refresh*, which is a file system list operation to pick up new files from the primary and possibly dropping its `MemStore`. This refresh will only be performed if the `MemStore` size of the biggest secondary region replica is at least `hbase.region.replica.storefile.refresh.memstore.multiplier` times bigger than the biggest `MemStore` of a primary replica. The default value for `hbase.region.replica.storefile.refresh.memstore.multiplier` is 4.



Note

If this operation is performed, the secondary replica might obtain partial row updates across column families because column families are flushed independently. Hortonworks recommends that you configure HBase to perform this operation infrequently.

You can disable this feature by setting the value to a large number, but this might cause replication to be blocked without resolution.

Secondary Replica Failover

When a secondary region replica first comes online, or after a secondary region fails over, it may have contain edits from its `MemStore`. The secondary replica must ensure that it accesses stale data that has been overwritten before serving requests after assignment. Therefore, the secondary waits until it detects a full flush cycle, consisting of start flush and commit flush, or a *region open event* replicated from the primary replica.

Until the flush cycle occurs, the secondary region replica rejects all read requests by way of an `IOException` with the following message:

The region's reads are disabled

Other replicas might be available to read, thus not causing any impact for the RPC with TIMELINE consistency.

To facilitate faster recovery, the secondary region triggers a flush request from the primary when it is opened. The configuration property `hbase.region.replica.wait.for.primary.flush`, which is enabled by default, can be used to disable this feature if needed.

2.4. Configuring HA Reads for HBase

To enable high availability for HBase reads, specify the following server-side and client-side configuration properties in your `hbase-site.xml` configuration file, and then restart the HBase Master and Region Servers.

The following table describes server-side properties. Set these properties for all servers in your HBase cluster that will use region replicas.

Property	Example value	Description
<code>hbase.regionserver.storefile.refresh.period</code>	30000	<p>Specifies the period (in milliseconds) for refreshing the store files for secondary regions. The default value is 0, which indicates that the feature is disabled. Secondary regions receive new files from the primary region after the secondary replica refreshes the list of files in the region.</p> <p>Note: Too-frequent refreshes might cause extra Namenode pressure. If files cannot be refreshed for longer than HFile TTL, specified with <code>hbase.master.hfilecleaner.ttl</code>, the requests are rejected.</p> <p>Refresh period should be a non-zero number if META replicas are enabled (see <code>hbase.meta.replica.count</code>).</p> <p>If you specify refresh period, we recommend configuring HFile TTL to a larger value than its default.</p>
<code>hbase.region.replica.replication.enabled</code>	true	<p>Determines whether asynchronous WAL replication is enabled or not. The value can be true or false. The default is false.</p> <p>If this property is enabled, a replication peer named <code>region_replica_replication</code> is created. The replication peer replicates changes to region replicas for any tables that have region replication set to 1 or more.</p> <p>After enabling this property, disabling it requires setting it to false and disabling the replication peer using the shell or the <code>ReplicationAdmin</code></p>

Property	Example value	Description
		java class. When replication is explicitly disabled and then re-enabled, you must set <code>hbase.replication</code> to true.
<code>hbase.master.hfilecleaner.ttl</code>	3600000	Specifies the period (in milliseconds) to keep store files in the archive folder before deleting them from the file system.
<code>hbase.master.loadbalancer.class</code>	<code>org.apache.hadoop.hbase.master.balancer.StochasticLoadBalancer</code>	Specifies the Java class used for balancing the load of all HBase clients. The default value is <code>org.apache.hadoop.hbase.master.balancer.StochasticLoadBalancer</code> , which is the only load balancer that supports reading data from Region Servers in secondary mode.
<code>hbase.meta.replica.count</code>	3	Region replication count for the meta regions. The default value is 1.
<code>hbase.regionserver.storefile.refresh.all</code>	false	Determines whether all store files will be refreshed, as opposed to just META tables. The default is true. Set this value to false when <code>hbase.region.replica.replication.enabled</code> is true. This should be true if meta replicas are enabled (via <code>hbase.meta.replica.count</code> set to greater than 1).
<code>hbase.region.replica.wait.for.primary.flush</code>	true	Specifies whether to wait for a full flush cycle from the primary before starting to serve data in a secondary replica. Disabling this feature might cause secondary replicas to read stale data when a region is transitioning to another Region Server.
<code>hbase.region.replica.storefile.refresh.memstore.multiplier</code>	4	Multiplier for a "store file refresh" operation for the secondary region replica. If a Region Server has memory pressure, the secondary region will refresh its store files if the MemStore size of the biggest secondary replica is bigger than this multiplier times than the MemStore size of the biggest primary replica. To disable this feature (not recommended), set this property to a large value.

The following table lists client-side properties. Set these properties for all clients (applications) and servers (in your HBase cluster) that will use region replicas.

Property	Example value	Description
<code>hbase.ipc.client.specificThreadForWriting</code>	true	Specifies whether to enable interruption of RPC threads at the client side. This is required for

Property	Example value	Description
		region replicas with fallback RPC's to secondary regions.
<code>hbase.client.primaryCallTimeout.get</code>	10000	Specifies the timeout (in microseconds) before secondary fallback RPC's are submitted for get requests with <code>Consistency.TIMELINE</code> to the secondary replicas of the regions. The default value is 10ms. Setting this to a smaller value increases the number of RPC's, but lowers 99th-percentile latencies.
<code>hbase.client.primaryCallTimeout.multiget</code>	10000	Specifies the timeout (in microseconds) before secondary fallback RPC's are submitted for multi-get requests (<code>HTable.get(List<Get>)</code>) with <code>Consistency.TIMELINE</code> to the secondary replicas of the regions. The default value is 10ms. Setting this to a smaller value increases the number of RPC's, but lowers 99th-percentile latencies.
<code>hbase.client.primaryCallTimeout.scan</code>	1000000	Specifies the timeout (in microseconds) before secondary fallback RPC's are submitted for scan requests with <code>Consistency.TIMELINE</code> to the secondary replicas of the regions. The default value is 1 second. Setting this to a smaller value increases the number of RPC's, but lowers 99th-percentile latencies.
<code>hbase.meta.replicas.use</code>	true	Specifies whether to use META table replicas or not. The default value is false.

2.5. Creating Highly-Available HBase Tables

HBase tables are not highly available by default. To enable high availability, designate a table as HA during table creation.

Creating HA Tables with the HBase Java API

HBase application developers create highly available HBase tables programmatically using the Java API, as shown in the following example:

```
HTableDescriptor htd =
    new HTableDescriptor(TableName.valueOf("test_table"));
htd.setRegionReplication(2);
...
admin.createTable(htd);
```

This example creates a table named `test_table` that is replicated to one secondary region. To replicate `test_table` to two secondary replicas, pass 3 as a parameter to the `setRegionReplication()` method.

Creating HA Tables with the HBase Shell

Create HA tables using the `REGION_REPLICATION` keyword with the HBase shell. Valid values are 1, 2, and 3, indicating the total number of copies. The default value is 1.

The following example creates a table named `t1` that is replicated to one secondary replica:

```
CREATE 't1', 'f1', {REGION_REPLICATION => 2}
```

To replicate `t1` to two secondary regions, set `REGION_REPLICATION` to 3:

```
CREATE 't1', 'f1', {REGION_REPLICATION => 3}
```

2.6. Querying Secondary Regions

This section describes how to query HA-enabled HBase tables.

Querying HBase with the Java API

The HBase Java API enables application developers to specify the desired data consistency for a query using the `setConsistency()` method, as shown in the following example. A new enum, `CONSISTENCY`, specifies two levels of data consistency: `TIMELINE` and `STRONG`.

```
Get get = new Get(row);
get.setConsistency(CONSISTENCY.TIMELINE);
...
Result result = table.get(get);
```

HBase application developers can also pass multiple `gets`:

```
Get get1 = new Get(row);
get1.setConsistency(Consistency.TIMELINE);
...
ArrayList<Get> gets = new ArrayList<Get>();
...
Result[] results = table.get(gets);
```

The `setConsistency()` method is also available for `Scan` objects:

```
Scan scan = new Scan();
scan.setConsistency(CONSISTENCY.TIMELINE);
...
ResultScanner scanner = table.getScanner(scan);
```

In addition, you can use the `Result.isStale()` method to determine whether the query results arrived from the primary or a secondary replica:

```
Result result = table.get(get);
if (result.isStale()) {
    ...
}
```

Querying HBase Interactively

To specify the desired data consistency for each query, use the HBase shell:

```
hbase(main):001:0> get 't1', 'r6', {CONSISTENCY => "TIMELINE"}
```

Interactive scans also accept this syntax:

```
hbase(main):001:0> scan 't1', {CONSISTENCY => 'TIMELINE'}
```



Note

This release of HBase does not provide a mechanism to determine if the results from an interactive query arrived from the primary or a secondary replica.

You can also request a specific region replica for debugging:

```
hbase> get 't1', 'r6', {REGION_REPLICA_ID=>0, CONSISTENCY=>'TIMELINE'}
hbase> get 't1', 'r6', {REGION_REPLICA_ID=>2, CONSISTENCY=>'TIMELINE'}
```

2.7. Monitoring Secondary Region Replicas

HBase provides highly available tables by replicating table regions. All replicated regions have a unique replica ID. The replica ID for a primary region is always 0. The HBase web-based interface displays the replica IDs for all defined table regions. In the following example, the table `t1` has two regions. The secondary region is identified by a replica ID of 1.

Table `t1`

Table Attributes

Attribute Name	Value	Description
Enabled	true	Is the table enabled
Compaction	NONE	Is the table compacting

Table Regions

Name	Region Server	Start Key	End Key	Requests	ReplicaID
t1_1399592796208.b98c5ea1129dd91cf73a28fa625ea67c.	sandbox.hortonworks.com:60020			0	0
t1_1399592796208_0001.567fbacba4257396c3aaa07b2fcf5032.	sandbox.hortonworks.com:60020			0	1

Regions by Region Server

Region Server	Region Count
sandbox.hortonworks.com:60020	2

To access the HBase Master Server user interface, point your browser to port 60010.

3. Namenode High Availability

The HDFS NameNode High Availability feature enables you to run redundant NameNodes in the same cluster in an Active/Passive configuration with a hot standby. This eliminates the NameNode as a potential single point of failure (SPOF) in an HDFS cluster.

Before the release of HDFS NameNode High Availability, if a cluster had a single NameNode, and that machine or process became unavailable, the entire cluster would be unavailable until the NameNode was either restarted or started on a separate machine. This situation impacted the total availability of the HDFS cluster in two major ways:

- In the case of an unplanned event such as a server failure, the cluster would be unavailable until an operator restarted the NameNode.
- Planned maintenance events such as software or hardware upgrades on the NameNode machine would result in periods of cluster downtime.

HDFS NameNode HA avoids this by facilitating either a fast failover to the new NameNode during machine crash, or a graceful administrator-initiated failover during planned maintenance.

This guide provides an overview of the HDFS NameNode High Availability (HA) feature, instructions on how to deploy Hue with an HA cluster, and instructions on how to enable HA on top of an existing HDP cluster using the Quorum Journal Manager (QJM) and Zookeeper Failover Controller for configuration and management. Using the QJM and Zookeeper Failover Controller enables the sharing of edit logs between the Active and Standby NameNodes.



Note

This guide assumes that an existing HDP cluster has been manually installed and deployed. If your existing HDP cluster was installed using Ambari, configure NameNode HA using the Ambari wizard, as described in "Enabling NameNode High Availability" in the [Ambari User's Guide](#).

3.1. Architecture

In a typical HA cluster, two separate machines are configured as NameNodes. In a working cluster, one of the NameNode machine is in the **Active** state, and the other is in the **Standby** state.

The Active NameNode is responsible for all client operations in the cluster, while the Standby acts as a slave. The Standby machine maintains enough state to provide a fast failover (if required).

In order for the Standby node to keep its state synchronized with the Active node, both nodes communicate with a group of separate daemons called **JournalNodes (JNs)**. When the Active node performs any namespace modification, the Active node durably logs a modification record to a majority of these JNs. The Standby node reads the edits from the JNs and continuously watches the JNs for changes to the edit log. Once the Standby Node observes the edits, it applies these edits to its own namespace. When using QJM,


JournalNodes acts the shared editlog storage. In a failover event, the Standby ensures that it has read all of the edits from the JournalNodes before promoting itself to the Active state. (This mechanism ensures that the namespace state is fully synchronized before a failover completes.)



Note

Secondary NameNode is not required in HA configuration because the Standby node also performs the tasks of the Secondary NameNode.

To provide a fast failover, it is also necessary that the Standby node have up-to-date information on the location of blocks in your cluster. To get accurate information about the block locations, DataNodes are configured with the location of both of the NameNodes, and send block location information and heartbeats to both NameNode machines.



It is vital for the correct operation of an HA cluster that only one of the NameNodes should be active at a time. Failure to do so, would cause the namespace state to quickly diverge between the two NameNode machines thus causing potential data loss. This situation is called a *split-brain scenario*.

To prevent the split-brain scenario, the JournalNodes allow only one NameNode to be a writer at a time. During failover, the NameNode that is to be chosen to become active takes over the role of writing to the JournalNodes. This process prevents the other NameNode from continuing in the active state and thus lets the new active node proceed with the failover safely.

3.2. Hardware Resources

Ensure that you prepare the following hardware resources:

- **NameNode machines:** The machines where you run Active and Standby NameNodes, should have exactly the same hardware. For recommended hardware for NameNodes, see "Hardware for Master Nodes" in the [Cluster Planning Guide](#).

- **JournalNode machines:** The machines where you run the JournalNodes. The JournalNode daemon is relatively lightweight, so these daemons may reasonably be co-located on machines with other Hadoop daemons, for example the NameNodes or the YARN ResourceManager.



Note

There must be at least three JournalNode daemons, because edit log modifications must be written to a majority of JournalNodes. This lets the system tolerate failure of a single machine. You may also run more than three JournalNodes, but in order to increase the number of failures that the system can tolerate, you must run an odd number of JournalNodes. For example, 3, 5, 7, and so on.

When you run $\langle N \rangle$ number of JournalNodes, the system can tolerate at most $(\langle N \rangle - 1) / 2$ failures and continue to function normally.

- **Zookeeper machines:** For automated failover functionality, there must be an existing Zookeeper cluster available. The Zookeeper service nodes can be co-located with other Hadoop daemons.

In an HA cluster, the Standby NameNode also performs checkpoints of the namespace state. Therefore, do not deploy a Secondary NameNode, CheckpointNode, or BackupNode in an HA cluster.

3.3. Deploy NameNode HA Cluster

HA configuration is backward compatible and works with your existing single NameNode configuration. The following instructions describe how to set up NameNode HA on a manually-installed cluster. If you installed with Ambari and manage HDP on Ambari 1.7.0 or later, instead of the manual instructions use the Ambari documentation, "Enabling NameNode High Availability" in the [Ambari User's Guide](#).



Note

HA cannot accept HDFS cluster names that include an underscore (`_`).

To deploy a NameNode HA cluster, use the steps in the following subsections.

3.3.1. 1. Configure NameNode HA Cluster

First, add high availability configurations to your HDFS configuration files. Start by taking the HDFS configuration files from the original NameNode in your HDP cluster, and use that as the base, adding the properties mentioned below to those files.

After you have added the configurations below, ensure that the same set of HDFS configuration files are propagated to all nodes in the HDP cluster. This ensures that all the nodes and services are able to interact with the highly available NameNode.

Add the following configuration options to your `hdfs-site.xml` file:

- **dfs.nameservices**

Choose an arbitrary but logical name (for example mycluster) as the value for `dfs.nameservices` option. This name will be used for both configuration and authority component of absolute HDFS paths in the cluster.

```
<property>
  <name>dfs.nameservices</name>
  <value>mycluster</value>
  <description>Logical name for this new nameservice</description>
</property>
```

If you are also using HDFS Federation, this configuration setting should also include the list of other nameservices, HA or otherwise, as a comma-separated list.

- **dfs.ha.namenodes.[*\$nameservice ID*]**

Provide a list of comma-separated NameNode IDs. DataNodes use this this property to determine all the NameNodes in the cluster.

For example, for the nameservice ID `mycluster` and individual NameNode IDs `nn1` and `nn2`, the value of this property is:

```
<property>
  <name>dfs.ha.namenodes.mycluster</name>
  <value>nn1,nn2</value>
  <description>Unique identifiers for each NameNode in the nameservice</description>
</property>
```



Note

Currently, a maximum of two NameNodes can be configured per nameservice.

- **dfs.namenode.rpc-address.[*\$nameservice ID*].[*\$name node ID*]**

Use this property to specify the fully-qualified RPC address for each NameNode to listen on.

Continuing with the previous example, set the full address and IPC port of the NameNode process for the above two NameNode IDs - `nn1` and `nn2`.

Note that there will be two separate configuration options:

```
<property>
  <name>dfs.namenode.rpc-address.mycluster.nn1</name>
  <value>machine1.example.com:8020</value>
</property>

<property>
  <name>dfs.namenode.rpc-address.mycluster.nn2</name>
  <value>machine2.example.com:8020</value>
</property>
```

- **dfs.namenode.http-address.[*\$nameservice ID*].[*\$name node ID*]**

Use this property to specify the fully-qualified HTTP address for each NameNode to listen on.

Set the addresses for both NameNodes HTTP servers to listen on. For example:

```
<property>
  <name>dfs.namenode.http-address.mycluster.nn1</name>
  <value>machine1.example.com:50070</value>
</property>

<property>
  <name>dfs.namenode.http-address.mycluster.nn2</name>
  <value>machine2.example.com:50070</value>
</property>
```



Note

If you have Hadoop security features enabled, set the https-address for each NameNode.

- **dfs.namenode.shared.edits.dir**

Use this property to specify the URI that identifies a group of JournalNodes (JNs) where the NameNode will write and read edits.

Configure the addresses of the JNs that provide the shared edits storage. The Active nameNode writes to this shared storage and the Standby NameNode reads from this location to stay up-to-date with all the file system changes.

Although you must specify several JournalNode addresses, you must configure only one of these URIs for your cluster.

The URI should be of the form:

```
qjournal://host1:port1;host2:port2;host3:port3/journalId
```

The Journal ID is a unique identifier for this nameservice, which allows a single set of JournalNodes to provide storage for multiple federated namesystems. You can reuse the nameservice ID for the journal identifier.

For example, if the JournalNodes for a cluster were running on `node1.example.com`, `node2.example.com` and `node3.example.com`, and the nameservice ID were `mycluster`, you would use the following value for this setting:

```
<property>
  <name>dfs.namenode.shared.edits.dir</name>
  <value>qjournal://node1.example.com:8485;node2.example.com:8485;node3.
example.com:8485/mycluster</value>
</property>
```



Note

Note that the default port for the JournalNode is 8485.

- **dfs.client.failover.proxy.provider.[Nameservice ID]**

This property specifies the Java class that HDFS clients use to contact the Active NameNode. DFS Client uses this Java class to determine which NameNode is the current Active and therefore which NameNode is currently serving client requests.

Use the `ConfiguredFailoverProxyProvider` implementation if you are not using a custom implementation.

For example:

```
<property>
  <name>dfs.client.failover.proxy.provider.mycluster</name>
  <value>org.apache.hadoop.hdfs.server.namenode.ha.
ConfiguredFailoverProxyProvider</value>
</property>
```

- **dfs.ha.fencing.methods**

This property specifies a list of scripts or Java classes that will be used to fence the Active NameNode during a failover.

For maintaining system correctness, it is important to have only one NameNode in the Active state at any given time. When using the Quorum Journal Manager, only one NameNode will ever be allowed to write to the JournalNodes, so there is no potential for corrupting the file system metadata from a split-brain scenario. However, when a failover occurs, it is still possible that the previous Active NameNode could serve read requests to clients, which may be out of date until that NameNode shuts down when trying to write to the JournalNodes.

For this reason, it is still recommended to configure some fencing methods even when using the Quorum Journal Manager. To improve the availability of the system in the event the fencing mechanisms fail, it is advisable to configure a fencing method which is guaranteed to return success as the last fencing method in the list. Note that if you choose to use no actual fencing methods, you must set some value for this setting, for example `shell(/bin/true)`.

The fencing methods used during a failover are configured as a carriage-return-separated list, which will be attempted in order until one indicates that fencing has succeeded. The following two methods are packaged with Hadoop: `shell` and `sshfence`. For information on implementing a custom fencing method, see the `org.apache.hadoop.ha.NodeFencer` class.

- **sshfence**: connects to the Active NameNode using the Secure SHell (SSH) and kill the process.

The `sshfence` option uses SSH to connect to the target node and uses `fuser` to kill the process listening on the TCP port of the service. In order for this fencing option to work, it must be able to use SSH to connect to the target node without providing a passphrase. Ensure that you configure the `dfs.ha.fencing.ssh.private-key-files` option, which is a comma-separated list of SSH private key files.

For example:

```

<property>
  <name>dfs.ha.fencing.methods</name>
  <value>sshfence</value>
</property>

<property>
  <name>dfs.ha.fencing.ssh.private-key-files</name>
  <value>/home/exampleuser/.ssh/id_rsa</value>
</property>

```

Optionally, you can also configure a non-standard username or port to perform the SSH. You can also configure a timeout, in milliseconds, for the SSH, after which this fencing method fails. To configure a non-standard username or port and timeout, see the properties given below:

```

<property>
  <name>dfs.ha.fencing.methods</name>
  <value>sshfence([[username][:port]])</value>
</property>

<property>
  <name>dfs.ha.fencing.ssh.connect-timeout</name>
  <value>30000</value>
</property>

```

- **shell:** Run an arbitrary shell command to fence the Active NameNode.

The shell fencing method runs an arbitrary shell command:

```

<property>
  <name>dfs.ha.fencing.methods</name>
  <value>shell(/path/to/my/script.sh arg1 arg2 ...)</value>
</property>

```

The string between the parentheses [(and)] is passed directly to a bash shell and may not include any closing parentheses.

The shell command will be run with an environment set up to contain all of the current Hadoop configuration variables, with the underscore character (_) replacing any dot or period characters (.) in the configuration keys. The configuration used has already had any namenode-specific configurations promoted to their generic forms. For example, `dfs_namenode_rpc-address` contains the RPC address of the target node, even though the configuration may specify that variable as `dfs.namenode.rpc-address.ns1.nn1`.

Additionally, the following variables that refer to the target node to be fenced are also available:

- `$target_host`: Hostname of the node to be fenced
- `$target_port`: IPC port of the node to be fenced
- `$target_address`: The combination of `$target_host` and `$target_port` as `host:port`
- `$target_nameserviceid`: The nameservice ID of the NN to be fenced

- `$target_namenodeid`: The namenode ID of the NN to be fenced

These environment variables may also be used as substitutions in the shell command. For example:

```
<property>
  <name>dfs.ha.fencing.methods</name>
  <value>shell(/path/to/my/script.sh --nameservice=
$target_nameserviceid $target_host:$target_port)</value>
</property>
```

If the shell command returns an exit code of 0, the fencing is successful.



Note

This fencing method does not implement any timeout. If timeouts are necessary, they should be implemented in the shell script itself. For example, by forking a subshell to kill its parent in a specified number of seconds.

- **fs.defaultFS** The default path prefix used by the Hadoop FS client. Optionally, you may now configure the default path for Hadoop clients to use the new HA-enabled logical URI. For example, for mycluster nameservice ID, this will be the value of the authority portion of all of your HDFS paths. Configure this property in the `core-site.xml` file:

```
<property>
  <name>fs.defaultFS</name>
  <value>hdfs://mycluster</value>
</property>
```

- **dfs.journalnode.edits.dir** This is the absolute path on the JournalNode servers where the edits and other local state information used by the JournalNodes are stored. You may only use a single path for this configuration. Redundancy for this data is provided by either running multiple separate JournalNodes or by configuring this directory on a locally-attached RAID array. For example:

```
<property>
  <name>dfs.journalnode.edits.dir</name>
  <value>/path/to/journal/node/local/data</value>
</property>
```



Note

See "Creating Service Principals and Keytab files for HDP" in [Installing HDP Manually](#) for instructions on configuring Kerberos-based security for Highly Available clusters.

3.3.2. 2. Deploy NameNode HA Cluster

In this section, we use NN1 to denote the original NameNode in the non-HA setup, and NN2 to denote the other NameNode that is to be added in the HA setup.



Note

HA clusters reuse the nameservice ID to identify a single HDFS instance that may consist of multiple HA NameNodes.

A new abstraction called NameNode ID is added with HA. Each NameNode in the cluster has a distinct NameNode ID to distinguish it.

To support a single configuration file for all of the NameNodes, the relevant configuration parameters are suffixed with both the nameservice ID and the NameNode ID.

1. Start the JournalNode daemons on the set of machines where the JournalNodes are deployed. On each machine, execute the following command:

```
su -l hdfs -c "/usr/hdp/current/hadoop-hdfs-journalnode/./hadoop/sbin/hadoop-daemon.sh start journalnode"
```

2. Wait for the daemon to start on each of the JournalNode machines.

3. Initialize JournalNodes.

- At the NN1 host machine, execute the following command:

```
su -l hdfs -c "namenode -initializeSharedEdits -force"
```

This command formats all the JournalNodes. This by default happens in an interactive way. The command prompts users for "Y/N" input to confirm the format. You can skip the prompt by using option **-force** or **-nonInteractive**.

It also copies all the edits data after the most recent checkpoint from the edits directories of the local NameNode (NN1) to JournalNodes.

- At the host with the journal node (if it is separated from the primary host), execute the following command:

```
su -l hdfs -c "namenode -initializeSharedEdits -force"
```

- Initialize HA state in ZooKeeper. Execute the following command on NN1:

```
hdfs zkfc -formatZK -force
```

This command creates a znode in ZooKeeper. The failover system stores uses this znode for data storage.

- Check to see if Zookeeper is running. If not, start Zookeeper by executing the following command on the ZooKeeper host machine(s).

```
su - zookeeper -c "export ZOOCFGDIR=/usr/hdp/current/zookeeper-server/conf ; export ZOOCFG=zoo.cfg;
                    source /usr/hdp/current/zookeeper-server/conf/
zookeeper-env.sh ; /usr/hdp/current/zookeeper-server/bin/zkServer.sh
start"
```

- At the standby namenode host, execute the following command:

```
su -l hdfs -c "namenode -bootstrapStandby -force"
```

4. Start NN1. At the NN1 host machine, execute the following command:

```
su -l hdfs -c "/usr/hdp/current/hadoop-hdfs-namenode/./hadoop/sbin/hadoop-daemon.sh start namenode"
```

Make sure that NN1 is running correctly.

5. Format NN2 and copy the latest checkpoint (FSImage) from NN1 to NN2 by executing the following command:

```
su -l hdfs -c "namenode -bootstrapStandby -force"
```

This command connects with NN1 to get the namespace metadata and the checkpointed fsimage. This command also ensures that NN2 receives sufficient editlogs from the JournalNodes (corresponding to the fsimage). This command fails if JournalNodes are not correctly initialized and cannot provide the required editlogs.

6. Start NN2. Execute the following command on the NN2 host machine:

```
su -l hdfs -c "/usr/hdp/current/hadoop-hdfs-namenode/./hadoop/sbin/hadoop-daemon.sh start namenode"
```

Ensure that NN2 is running correctly.

7. Start DataNodes. Execute the following command on all the DataNodes:

```
su -l hdfs -c "/usr/hdp/current/hadoop-hdfs-datanode/./hadoop/sbin/hadoop-daemon.sh start datanode"
```

8. Validate the HA configuration.

Go to the NameNode web pages separately by browsing to their configured HTTP addresses. Under the configured address label, you should see the HA state of the NameNode. The NameNode can be either in "standby" or "active" state.

NameNode 'example.com:8020' (standby)

Started:	Thu Aug 15 02:16:35 UTC 2013
Version:	3.0.0-SNAPSHOT, 5c35d30ce6f27a7d452e398be48be3f0a403e286
Compiled:	2013-08-14T19:42Z by hdfs from trunk
Cluster ID:	CID-9165ed44-7149-4598-a4a5-6259f5d12689
Block Pool ID:	BP-2092817692-68.142.245.166-1375143516059

[NameNode Logs](#)



Note

The HA NameNode is initially in the Standby state after it is bootstrapped. You can use either the Java Management Extension (JMX) tag `.HAState` to query the HA state of a NameNode or the following command:

```
hdfs haadmin -getServiceState
```

9. Transition one of the HA NameNodes to Active state.

Initially, both NN1 and NN2 are in Standby state. Therefore you must transition one of the NameNodes to Active state. This transition can be performed using one of the following options:

- **Option I - Using CLI** Use the command line interface (CLI) to transition one of the NameNodes to Active State. Execute the following command on that NameNode host server:

```
hdfs haadmin -failover --forcefence --forceactive <serviceId> <namenodeId>
```

For more information on the `haadmin` command, see "[Appendix: Administrative Commands](#)."

- **Option II - Deploying Automatic Failover** You can configure and deploy automatic failover using the instructions provided in [Configure and Deploy NameNode Automatic Failover](#).

3.3.3.3. Deploy Hue with an HA Cluster

If you are going to use Hue with an HA Cluster, make the following changes to `/etc/hue/conf/hue.ini`:

1. Install the Hadoop HttpFS component on the Hue server.

For RHEL/CentOS/Oracle Linux:

```
yum install hadoop-httpfs
```

For SLES:

```
yum install hadoop-httpfs
```

2. Modify `/etc/hadoop-httpfs/conf/httpfs-env.sh` to add the JDK path. In the file, ensure that `JAVA_HOME` is set:

```
export JAVA_HOME=/usr/jdk64/jdk1.7.0_67
```

3. Configure the HttpFS service script for use by setting up the symlink in `/etc/init.d`:

```
> ln -s /usr/hdp/{HDP2.2.x version number}/etc/rc.d/init.d/hadoop-httpfs
    /etc/init.d/hadoop-httpfs
```

For example, `{HDP2.2.x version number}` could be `2.2.0.0-2041`.

4. Modify `/etc/hadoop-httpfs/conf/httpfs-site.xml` to configure HttpFS to talk to the cluster, by confirming that the following properties are correct:

```
<property>
  <name>httpfs.proxyuser.hue.hosts</name>
  <value>*</value>
</property>

<property>
  <name>httpfs.proxyuser.hue.groups</name>
  <value>*</value>
</property>
```

5. Start the HttpFS service.

```
service hadoop-httpfs start
```

6. Modify the `core-site.xml` file. On the NameNodes and all the DataNodes, add the following properties to the `$HADOOP_CONF_DIR /core-site.xml` file, where `$HADOOP_CONF_DIR` is the directory for storing the Hadoop configuration files. For example, `/etc/hadoop/conf`.

```
<property>
  <name>hadoop.proxyuser.httpfs.groups</name>
  <value>*</value>
</property>

<property>
  <name>hadoop.proxyuser.httpfs.hosts</name>
  <value>*</value>
</property>
```

7. In the `hue.ini` file, under the `[hadoop][hdfs_clusters][[default]]` subsection, use the following variables to configure the cluster:

Property	Description	Example
<code>fs_defaults</code>	NameNode URL using the logical name for the new name service. For reference, this is the <code>dfs.nameservices</code> property in <code>hdfs-site.xml</code> in your Hadoop configuration.	<code>hdfs://mycluster</code>
<code>webhdfs_url</code>	URL to the HttpFS server.	<code>http://c6401.apache.org:14000/webhdfs/v1/</code>

8. Restart Hue for the changes to take effect.

```
service hue restart
```

3.3.4. 4. Deploy Oozie with an HA Cluster

You can configure multiple Oozie servers against the same database to provide High Availability (HA) for the Oozie service. You must meet the following prerequisites:

- A database that supports multiple concurrent connections. In order to have full HA, the database should also have HA support, or it becomes a single point of failure.



Note

The default Derby database does not support this.

- A ZooKeeper ensemble. Apache ZooKeeper is a distributed, open-source coordination service for distributed applications. Oozie servers use it for coordinating access to the database and communicating with each other. In order to have full HA, there should be at least 3 ZooKeeper servers. Find more information about Zookeeper [here](#).
- Multiple Oozie servers.



Important

While not strictly required, you should configure all ZooKeeper servers to have identical properties.

- A Loadbalancer, Virtual IP, or Round-Robin DNS. This is used to provide a single entry-point for users and for callbacks from the JobTracker. The load balancer should be configured for round-robin between the Oozie servers to distribute the requests. Users using either the Oozie client, a web browser, or the REST API should connect through the load balancer. In order to have full HA, the load balancer must also have HA support or it becomes a single point of failure. For information about how to set up your Oozie servers to handle failover, see [Configuring Oozie Failover](#).

3.4. Operating a NameNode HA Cluster

- While operating an HA cluster, the Active NameNode cannot commit a transaction if it cannot write successfully to a quorum of the JournalNodes.
- When restarting an HA cluster, the steps for initializing JournalNodes and NN2 can be skipped.
- Start the services in the following order:

1. JournalNodes
2. NameNodes



Note

Verify that the ZKFailoverController (ZKFC) process on each node is running so that one of the NameNodes can be converted to active state.

3. DataNodes

- In a NameNode HA cluster, the following `dfs admin` command options will run only on the active NameNode:

```
-rollEdits  
-setQuota  
-clrQuota  
-setSpaceQuota  
-clrSpaceQuota  
-setStoragePolicy  
-getStoragePolicy  
-finalizeUpgrade  
-rollingUpgrade  
-printTopology  
-allowSnapshot <snapshotDir>  
-disallowSnapshot <snapshotDir>
```

The following `dfs admin` command options will run on both the active and standby NameNodes:

```
-safemode enter
-saveNamespace
-restoreFailedStorage
-refreshNodes
-refreshServiceAcl
-refreshUserToGroupsMappings
-refreshSuperUserGroupsConfiguration
-refreshCallQueue
-metasave
-setBalancerBandwidth
```

The `-refresh <host:ipc_port> <key> arg1..argn` command will be sent to the corresponding host according to its command arguments.

The `-fetchImage <local directory>` command attempts to identify the active NameNode through an RPC call, and then fetch the `fsimage` from that NameNode. This means that usually the `fsimage` is retrieved from the active NameNode, but it is not guaranteed because a failover can happen between the two operations.

The following `dfs admin` command options are sent to the DataNodes:

```
-refreshNamenodes
-deleteBlockPool
-shutdownDatanode <datanode_host:ipc_port> upgrade
-getDatanodeInfo <datanode_host:ipc_port>
```

3.5. Configure and Deploy NameNode Automatic Failover

The preceding sections describe how to configure manual failover. In that mode, the system will not automatically trigger a failover from the active to the standby NameNode, even if the active node has failed. This section describes how to configure and deploy automatic failover.


Automatic failover adds the following components to an HDFS deployment:

- ZooKeeper quorum
- ZKFailoverController process (abbreviated as ZKFC).

The ZKFailoverController (ZKFC) is a ZooKeeper client that monitors and manages the state of the NameNode. Each of the servers that run the NameNode service also run a ZKFC. The ZKFC is responsible for:

- **Health monitoring:** The ZKFC periodically pings its local NameNode with a health-check command.
- **ZooKeeper session management:** When the local NameNode is healthy, the ZKFC holds a session open in ZooKeeper. If the local NameNode is active, it also holds a special "lock" znode. This lock uses ZooKeeper support for *ephemeral nodes*. If the session expires, the lock node is automatically deleted.
- **ZooKeeper-based election:** If the local NameNode is healthy and no other node currently holds the lock znode, the ZKFC attempts to acquire the lock. If the ZKFC succeeds,

then it has *won the election* and is responsible for running a failover to make its local NameNode active. The failover process is similar to the manual failover described above. First, the previous active node is fenced if necessary and then the local NameNode transitions to active state.



3.5.1. Prerequisites

Complete the following prerequisite configurations:

- Make sure that you have a working Zookeeper service. If you have an Ambari deployed HDP cluster with Zookeeper, you can use that. If not, deploy ZooKeeper using the instructions provided in the "Installing ZooKeeper" chapter of the [Installing HDP Manually](#) guide.



Note

In a typical deployment, ZooKeeper daemons are configured to run on three or five nodes. However, it is acceptable to co-locate the ZooKeeper nodes on the same hardware as the HDFS NameNode and Standby Node. Many operators choose to deploy the third ZooKeeper process on the same node as the YARN ResourceManager. To achieve performance and improve isolation, Hortonworks recommends configuring the ZooKeeper nodes such that the ZooKeeper data and HDFS metadata is stored on separate disk drives.

- Shut down your HA cluster that is configured for manual failover using the instructions provided in "Controlling HDP Services Manually," in the [HDP Reference Guide](#).

Currently, you cannot transition from a manual failover setup to an automatic failover setup while the cluster is running.

3.5.2. Instructions

Complete the following actions:

1. Configure automatic failover.

- a. Set up your cluster for automatic failover. Add the following property to the the `hdfs-site.xml` file for both the NameNode nodes:

```
<property>
  <name>dfs.ha.automatic-failover.enabled</name>
  <value>true</value>
</property>
```

- b. List the host-port pairs running the ZooKeeper service. Add the following property to the the `core-site.xml` file for both the NameNode machines:

```
<property>
  <name>ha.zookeeper.quorum</name>
  <value>zk1.example.com:2181,zk2.example.com:2181,zk3.example.
com:2181</value>
</property>
```



Note

Suffix the configuration key with the nameservice ID to configure the above settings on a per-nameservice basis. For example, in a cluster with federation enabled, you can explicitly enable automatic failover for only one of the nameservices by setting `dfs.ha.automatic-failover.enabled.$my-nameservice-id`.

2. Initialize HA state in ZooKeeper.

Execute the following command on NN1:

```
hdfs zkfc -formatZK -force
```

This command creates a znode in ZooKeeper. The automatic failover system stores uses this znode for data storage.

3. Check to see if Zookeeper is running. If not, start Zookeeper by executing the following command on the ZooKeeper host nodes.

```
su - zookeeper -c "export ZOOCFGDIR=/usr/hdp/current/zookeeper-server/
conf ;
export ZOOCFG=zoo.cfg; source /usr/hdp/current/zookeeper-server/conf/
zookeeper-env.sh ;
/usr/hdp/current/zookeeper-server/bin/zkServer.sh start"
```

4. Start the JournalNodes, NameNodes, and DataNodes using the instructions provided in "Controlling HDP Services Manually," in the [HDP Reference Guide](#).

5. Start the Zookeeper Failover Controller (ZKFC) by executing the following command:

```
su -l hdfs -c "/usr/hdp/current/hadoop-hdfs-namenode/./hadoop/sbin/hadoop-
daemon.sh start zkfc"
```

The sequence of starting the ZKFC determines which NameNode will become Active. For example, if the ZKFC is started on NN1 first, it causes NN1 to become Active.



Note

To convert a non-HA cluster to an HA cluster, Hortonworks recommends that you run the `bootstrapStandby` command to initialize NN2 before you start the ZKFC on any of the NameNode nodes.

6. Verify automatic failover.

a. Locate the Active NameNode.

Use the NameNode web UI to check the status for each NameNode host node.

b. Cause a failure on the Active NameNode host machine.

For example, you can use the following command to simulate a JVM failure:

```
kill -9 $PID_of_Active_NameNode
```

Or, you could power cycle the node or unplug its network interface to simulate outage.

c. The Standby NameNode should now automatically become Active within several seconds.



Note

The amount of time required to detect a failure and trigger a failover depends on the configuration of the `ha.zookeeper.session-timeout.ms` property. The default value for this property is 5 seconds.

d. If the test fails, your HA settings might be incorrectly configured.

Check the logs for the zkfc daemons and the NameNode daemons to diagnose the issue.

3.5.3. Configuring Oozie Failover

1. Set up your database for high availability. For details, see the documentation for your Oozie database.

Oozie database configuration properties may need special configuration. For details, see the JDBC driver documentation for your database.

2. Configure Oozie identically on two or more servers.

3. Set the `OOZIE_HTTP_HOSTNAME` variable in `oozie-env.sh` to the Load Balancer or Virtual IP address.

4. Start all Oozie servers.

5. Use either a Virtual IP Address or Load Balancer to direct traffic to Oozie servers.

6. Access Oozie via the Virtual IP or Load Balancer address.

3.6. Appendix: Administrative Commands

The subcommands of `hdfs haadmin` are used for administering an HA cluster.

Running the `hdfs haadmin` command without any additional arguments will display the following usage information:

```
Usage: DFSHAAdmin [-ns <nameserviceId>]
[-transitionToActive <serviceId>]
[-transitionToStandby <serviceId>]
[-failover [--forcefence] [--forceactive] <serviceId> <serviceId>]
[-getServiceState <serviceId>]
[-checkHealth <serviceId>]
[-help <command>]
```

This section provides high-level uses of each of these subcommands.

- **transitionToActive** and **transitionToStandby**: Transitions the state of the given NameNode to Active or Standby.

These subcommands cause a given NameNode to transition to the Active or Standby state, respectively. These commands do not attempt to perform any fencing, and thus should be used rarely. Instead, Hortonworks recommends using the following subcommand:

```
hdfs haadmin -failover
```

- **failover**: Initiates a failover between two NameNodes.

This subcommand causes a failover from the first provided NameNode to the second.

- If the first NameNode is in the Standby state, this command transitions the second to the Active state without error.
- If the first NameNode is in the Active state, an attempt is made to gracefully transition it to the Standby state. If this fails, the fencing methods that are configured by `dfs.ha.fencing.methods` are attempted in order until one succeeds. Only after this process can the second NameNode be transitioned to the Active state. If the fencing methods fail, the second NameNode is not transitioned to an Active state and an error is returned.
- **getServiceState**: Determines whether the given NameNode is Active or Standby.

This subcommand connects to the provided NameNode, determines its current state, and prints either "standby" or "active" to STDOUT appropriately. This subcommand might be used by cron jobs or monitoring scripts.

- **checkHealth**: Checks the health of the given NameNode.

This subcommand connects to the NameNode to check its health. The NameNode is capable of performing some diagnostics that include checking if internal services are running as expected. This command will return 0 if the NameNode is healthy else it will return a non-zero code.



Note

This subcommand is in implementation phase and currently always returns success unless the given NameNode is down.

4. Resource Manager High Availability

This chapter provides instructions on setting up the ResourceManager (RM) High Availability (HA) feature in an HDFS cluster. The Active and Standby ResourceManagers embed the Zookeeper-based ActiveStandbyElector to determine which ResourceManager should be active.



Note

This document assumes that an existing HDP cluster has been manually installed and deployed. It provides instructions on how to manually enable ResourceManager HA on top of the existing cluster.

The ResourceManager is a single point of failure (SPOF) in an HDFS cluster. Each cluster has a single ResourceManager, and if that machine or process becomes unavailable, the entire cluster is unavailable until the ResourceManager is either restarted or started on a separate node. This situation impacts the total availability of the HDFS cluster in two ways:

- Unplanned events, such as a node failure, cause the cluster to be unavailable until an operator restarts the ResourceManager.
- Planned maintenance events, such as software or hardware upgrades on the ResourceManager node, cause periods of cluster downtime.

The ResourceManager HA feature addresses these problems. This feature enables you to run redundant ResourceManagers in the same cluster in an Active/Passive configuration with a hot standby. This mechanism thus facilitates either a fast failover to the standby ResourceManager during node failure, or a graceful administrator-initiated failover during planned maintenance.

4.1. Hardware Resources

Ensure that you prepare the following hardware resources:

- **ResourceManager nodes:** The nodes where you run Active and Standby ResourceManagers should have exactly the same hardware. For recommended hardware for ResourceManagers, see "Hardware for Master Nodes" in the [Cluster Planning Guide](#).
- **Zookeeper machines:** For automated failover functionality, there must be an existing Zookeeper cluster available. The Zookeeper service nodes can be co-located with other Hadoop daemons.

4.2. Deploy ResourceManager HA Cluster

HA configuration is backward-compatible and works with your existing single ResourceManager configuration.

As described in the following sections, first configure manual or automatic ResourceManager failover. Then deploy the ResourceManager HA cluster.

4.2.1. Configure Manual or Automatic ResourceManager Failover

Prerequisites

Complete the following prerequisites:

- Make sure that you have a working Zookeeper service. If you had an Ambari- deployed HDP cluster with Zookeeper, you can use that Zookeeper service. If not, deploy ZooKeeper using the instructions provided in the "Installing ZooKeeper" chapter of the [Installing HDP Manually](#) guide.



Note

In a typical deployment, ZooKeeper daemons are configured to run on three or five nodes. However, it is acceptable to co-locate the ZooKeeper nodes on the same hardware as the HDFS NameNode and Standby Node. Many operators choose to deploy the third ZooKeeper process on the same node as the YARN ResourceManager. To achieve performance and improve isolation, Hortonworks recommends configuring the ZooKeeper nodes so that the ZooKeeper data and HDFS metadata are stored on separate disk drives.

- Shut down the cluster using the instructions provided in "Controlling HDP Services Manually," in the [HDP Reference Guide](#).

Set Common ResourceManager HA Properties

The following properties are required for both manual and automatic ResourceManager HA. Add these properties to the `etc/hadoop/conf/yarn-site.xml` file:

Property Name	Recommended Value	Description
<code>yarn.resourcemanager.ha.enabled</code>	<code>true</code>	Enable RM HA
<code>yarn.resourcemanager.ha.rm-ids</code>	Cluster-specific, e.g., <code>rm1,rm2</code>	A comma-separated list of ResourceManager IDs in the cluster.
<code>yarn.resourcemanager.hostname.<rm-id></code>	Cluster-specific	The host name of the ResourceManager. Must be set for all RMs.
<code>yarn.resourcemanager.recovery.enabled</code>	<code>true</code>	Enable job recovery on RM restart or failover.
<code>yarn.resourcemanager.store.class</code>	<code>org.apache.hadoop.yarn.server.resourcemanager.recovery.ZKRMStateStore</code>	The RMStateStore implementation to use to store the ResourceManager's internal state. The ZooKeeper- based store supports fencing implicitly, i.e., allows a single ResourceManager to make multiple changes at a time, and hence is recommended.
<code>yarn.resourcemanager.zk-address</code>	Cluster-specific	The ZooKeeper quorum to use to store the ResourceManager's internal state. For multiple ZK servers, use commas to separate multiple ZK servers.
<code>yarn.client.failover-proxy-provider</code>	<code>org.apache.hadoop.yarn.client.</code>	When HA is enabled, the class to be used by Clients, AMs and NMs to

Property Name	Recommended Value	Description
	ConfiguredRMFailoverProxyProvider	failover to the Active RM. It should extend org.apache.hadoop.yarn.client.RMFailoverProxyProvider This is an optional configuration. The default value is "org.apache.hadoop.yarn.client.ConfiguredRMFailoverProxyProvider"



Note

You can also set values for each of the following properties in `yarn-site.xml`:

```
yarn.resourcemanager.address.<rm#id>
yarn.resourcemanager.scheduler.address.<rm#id>
yarn.resourcemanager.admin.address.<rm#id>
yarn.resourcemanager.resource#tracker.address.<rm#id>
yarn.resourcemanager.webapp.address.<rm#id>
```

If these addresses are not explicitly set, each of these properties will use

```
yarn.resourcemanager.hostname.<rm-id>:default_port
```

such as `DEFAULT_RM_PORT`, `DEFAULT_RM_SCHEDULER_PORT`.

The following is a sample `yarn-site.xml` file with these common ResourceManager HA properties configured:

```
<!-- RM HA Configurations-->

<property>
  <name>yarn.resourcemanager.ha.enabled</name>
  <value>>true</value>
</property>

<property>
  <name>yarn.resourcemanager.ha.rm-ids</name>
  <value>rm1,rm2</value>
</property>

<property>
  <name>yarn.resourcemanager.hostname.rm1</name>
  <value>${rm1 address}</value>
</property>

<property>
  <name>yarn.resourcemanager.hostname.rm2</name>
  <value>${rm2 address}</value>
</property>

<property>
  <name>yarn.resourcemanager.webapp.address.rm1</name>
  <value>rm1_web_address:port_num</value>
  <description>We can set rm1_web_address separately. If not, it will use
  ${yarn.resourcemanager.hostname.rm1}:DEFAULT_RM_WEBAPP_PORT</
description>
```

```

</property>

<property>
  <name>yarn.resourcemanager.webapp.address.rm2</name>
  <value>rm2_web_address:port_num</value>
</property>

<property>
  <name>yarn.resourcemanager.recovery.enabled</name>
  <value>>true</value>
</property>

<property>
  <name>yarn.resourcemanager.store.class</name>
  <value>org.apache.hadoop.yarn.server.resourcemanager.recovery.
ZKRMStateStore</value>
</property>

<property>
  <name>yarn.resourcemanager.zk-address</name>
  <value>${zk1.address,zk2.address}</value>
</property>

<property>
  <name>yarn.client.failover-proxy-provider</name>
  <value>org.apache.hadoop.yarn.client.ConfiguredRMFailoverProxyProvider</
value>
</property>

```

Configure Manual ResourceManager Failover

Automatic ResourceManager failover is enabled by default, so it must be disabled for manual failover.

To configure manual failover for ResourceManager HA, add the `yarn.resourcemanager.ha.automatic-failover.enabled` configuration property to the `etc/hadoop/conf/yarn-site.xml` file, and set its value to "false":

```

<property>
  <name>yarn.resourcemanager.ha.automatic-failover.enabled</name>
  <value>>false</value>
</property>

```

Configure Automatic ResourceManager Failover

The preceding section described how to configure manual failover. In that mode, the system will not automatically trigger a failover from the active to the standby ResourceManager, even if the active node fails. This section describes how to configure automatic failover.

1. Add the following configuration options to the `yarn-site.xml` file:

Property Name	Recommended Value	Description
<code>yarn.resourcemanager.ha.automatic-failover.zk-base-path</code>	<code>/yarn-leader-election</code>	The base znode path to use for storing leader information, when using ZooKeeper-based leader election. This is an optional configuration. The default value is

Property Name	Recommended Value	Description
		/yarn-leader-election
yarn.resourcemanager.cluster-id	yarn-cluster	The name of the cluster. In a HA setting, this is used to ensure the RM participates in leader election for this cluster, and ensures that it does not affect other clusters.

Example:

```
<property>
  <name>yarn.resourcemanager.ha.automatic-failover.zk-base-path</name>
  <value>/yarn-leader-election</value>
  <description>Optional setting. The default value is /yarn-leader-
election</description>
</property>

<property>
  <name>yarn.resourcemanager.cluster-id</name>
  <value>yarn-cluster</value>
</property>
```

2. Automatic ResourceManager failover is enabled by default.

If you previously configured manual ResourceManager failover by setting the value of `yarn.resourcemanager.ha.automatic-failover.enabled` to **false**, you must delete this property to return automatic failover to its default enabled state.

4.2.2. Deploy the ResourceManager HA Cluster

1. Copy the `etc/hadoop/conf/yarn-site.xml` file from the primary ResourceManager host to the standby ResourceManager host.
2. Make sure that the `clientPort` value set in `etc/zookeeper/conf/zoo.cfg` matches the port set in the following `yarn-site.xml` property:

```
<property>
  <name>yarn.resourcemanager.zk-state-store.address</name>
  <value>localhost:2181</value>
</property>
```

3. Start ZooKeeper. Execute this command on the ZooKeeper host nodes:

```
su - zookeeper -c "export ZOOCFGDIR=/usr/hdp/current/zookeeper-server/
conf ;
export ZOOCFG=zoo.cfg; source /usr/hdp/current/zookeeper-server/conf/
zookeeper-env.sh ;
/usr/hdp/current/zookeeper-server/bin/zkServer.sh start "
```

4. Start HDFS using the instructions provided in "Controlling HDP Services Manually," in the [HDP Reference Guide](#).
5. Start YARN using the instructions provided in "Controlling HDP Services Manually," in the [HDP Reference Guide](#).
6. Set the active ResourceManager:

FOR MANUAL FAILOVER ONLY: If you configured manual ResourceManager failover, you must transition one of the ResourceManagers to Active mode. Execute the following CLI command to transition ResourceManager `rm1` to Active:

```
yarn rmadmin -transitionToActive rm1
```

You can use the following CLI command to transition ResourceManager `rm1` to Standby mode:

```
yarn rmadmin -transitionToStandby rm1
```

FOR AUTOMATIC FAILOVER: If you configured automatic ResourceManager failover, no action is required. The Active ResourceManager is chosen automatically.

7. Start all remaining unstarted cluster services using the instructions provided in the "Controlling HDP Services Manually" chapter of the [HDP Reference Guide](#).

4.2.3. Minimum Settings for Automatic ResourceManager HA Configuration

The minimum `yarn-site.xml` configuration settings for ResourceManager HA with automatic failover are as follows:

```
<property>
  <name>yarn.resourcemanager.ha.enabled</name>
  <value>>true</value>
</property>

<property>
  <name>yarn.resourcemanager.ha.rm-ids</name>
  <value>rm1,rm2</value>
</property>

<property>
  <name>yarn.resourcemanager.hostname.rm1</name>
  <value>192.168.1.9</value>
</property>

<property>
  <name>yarn.resourcemanager.hostname.rm2</name>
  <value>192.168.1.10</value>
</property>

<property>
  <name>yarn.resourcemanager.recovery.enabled</name>
  <value>>true</value>
</property>

<property>
  <name>yarn.resourcemanager.store.class</name>
  <value>org.apache.hadoop.yarn.server.resourcemanager.recovery.
ZKRMStateStore</value>
</property>

<property>
  <name>yarn.resourcemanager.zk-address</name>
```

```

    <value>192.168.1.9:2181,192.168.1.10:2181</value>
    <description>For multiple zk services, separate them with comma</
description>
</property>

<property>
  <name>yarn.resourcemanager.cluster-id</name>
  <value>yarn-cluster</value>
</property>

```

4.2.4. Testing ResourceManager HA on a Single Node

To test ResourceManager HA on a single node or launch more than one ResourceManager on a single node, you must add the following settings in the `yarn-site.xml` file.

To enable ResourceManager `rm1` to launch:

```

<property>
  <name>yarn.resourcemanager.ha.id</name>
  <value>rm1</value>
  <description>If we want to launch more than one RM in single node, we
  need this configuration</description>
</property>

```

To enable ResourceManager `rm2` to launch:

```

<property>
  <name>yarn.resourcemanager.ha.id</name>
  <value>rm2</value>
  <description>If we want to launch more than one RM in single node, we
  need this configuration</description>
</property>

```

You should also explicitly set values specific to each ResourceManager for the following properties separately in `yarn-site.xml`:

- `yarn.resourcemanager.address.<rm-id>`
- `yarn.resourcemanager.scheduler.address.<rm-id>`
- `yarn.resourcemanager.admin.address.<rm-id>`
- `yarn.resourcemanager.resource#tracker.address.<rm-id>`
- `yarn.resourcemanager.webapp.address.<rm-id>`

For example:

```

<!-- RM1 Configs -->

<property>
  <name>yarn.resourcemanager.address.rm1</name>
  <value>localhost:23140</value>
</property>

<property>
  <name>yarn.resourcemanager.scheduler.address.rm1</name>
  <value>localhost:23130</value>

```

```
</property>

<property>
  <name>yarn.resourcemanager.webapp.address.rm1</name>
  <value>localhost:23188</value>
</property>

<property>
  <name>yarn.resourcemanager.resource-tracker.address.rm1</name>
  <value>localhost:23125</value>
</property>

<property>
  <name>yarn.resourcemanager.admin.address.rm1</name>
  <value>localhost:23141</value>
</property>

<!-- RM2 configs -->
<property>
  <name>yarn.resourcemanager.address.rm2</name>
  <value>localhost:33140</value>
</property>

<property>
  <name>yarn.resourcemanager.scheduler.address.rm2</name>
  <value>localhost:33130</value>
</property>

<property>
  <name>yarn.resourcemanager.webapp.address.rm2</name>
  <value>localhost:33188</value>
</property>

<property>
  <name>yarn.resourcemanager.resource-tracker.address.rm2</name>
  <value>localhost:33125</value>
</property>

<property>
  <name>yarn.resourcemanager.admin.address.rm2</name>
  <value>localhost:33141</value>
</property>
```